

HIGH-PERFORMANCE SERVER SAN FOR WEB-SCALE DATA CENTERS

JOINT SOLUTION



USE CASES

HPC - Burst Buffer
M&E - Any-k video
Analytics

CHALLENGE

Sharing NVMe over the network with local performance

SOLUTION

NVmesh Server SAN

SPECS


- 4 hot swappable nodes in 2U
- 24 DIMMs per node
- 1 SIOM card support
- 2 low-profile PCI-E x16 slots

SOLIDSCALE REFERENCE ARCHITECTURE LEVERAGING SUPERMICRO BIGTWIN™, MELLANOX CONNECTX AND EXCELERO NVMesh®


New application workloads are driving the need for new storage architectures. Cloud and mobile applications, Industrial IoT and machine learning have a massive impact on the volumes of data that need to be transferred, stored and processed. Real-time analytics demand the lowest levels of latency. To meet these new requirements, companies need to rethink their storage strategies.



SCALABILITY



FLEXIBILITY



EFFICIENCY

Supermicro, Mellanox and Excelero have joined forces to deliver a high-performance, low-latency Server SAN solution leveraging SuperMicro BigTwin servers with NVMe flash, high-performance Mellanox ConnectX networking and Excelero NVMesh. The joint solution, which is based on the SolidScale Reference Architecture, is a Software-defined Block Storage platform that allows unmodified applications to utilize pooled NVMe storage devices across a network at local speeds and latencies. SolidScale was designed to combine virtually unlimited BigTwins to scale capacity and performance.

EXCELERO NVMesh

NVMesh is a Software-Defined Storage solution that enables customers to design Server SAN infrastructures for the most demanding enterprise and cloud-scale applications, leveraging standard servers and multiple tiers of flash. NVMesh features Elastic NVMe, a distributed block layer that allows any application to utilize pooled NVMe storage devices across a network at local speeds and latencies, while at the same time getting the benefits of centralized, redundant storage.

The primary benefit of NVMesh is that its patented Remote Direct Drive Access (RDDA) technology allows you to logically disaggregate NVMe drives in the compute nodes away from CPU resources. That is, though the local NVMe drive may be used by remote compute nodes, that usage does not consume local CPU. NVMesh is deployed as a virtual, distributed non-volatile array and supports both converged and disaggregated architectures, giving customers full freedom in their architectural design.

NVMesh KEY BENEFITS:

- Maximize the utilization of your NVMe SSD's.
- Predictable application performance – no noisy neighbors.
- Easy to manage & monitor, reduces the maintenance TCO.
- Mix different storage media types to optimize for cost, scale or performance.

SUPERMICRO BIGTWIN

The Supermicro BigTwin is the first and only 2U multi-node system supporting the highest performance processor, memory, storage, I/O and an incredible 30% better thermal capacity, leveraging the Titanium Level 96% efficient PowerStick power supplies (2200W/2600W). The 5th generation of Supermicro's Twin architecture is fully optimized for today's diverse workloads and architected for further technology breakthroughs.

The system delivers the highest performance and efficiency in a 2U 4-node platform. It supports the widest TDP range of CPUs (up to 205W), fully exploits all memory channels with 24 DIMMS per node and 24 All-Flash NVMe drives. Finally, the SuperMicro BigTwin provides 100% more I/O capacity and added flexibility with more than 10 networking options including 1GbE, 10G, 25G, 100G, IB and industry leading SIOM modular interconnect.

BIG TWIN BENEFITS:

- A Full Range of Processors Up to the highest performing 205W CPUs.
- Maximum Memory Full 24 DIMMs of memory per node.
- All-Flash NVMe 24 All-Flash NVMe or Hybrid Drives.
- Double the I/O Capacity Three PCI-E 3.0 x16 options per node.

MELLANOX END-TO-END ETHERNET SOLUTION

The Mellanox end-to-end Ethernet solution greatly simplifies the deployment and management of high-speed networking in the data center. Optimized for ultra-low latency lossless fabrics, the Mellanox Ethernet switches support the requirements of today's performance-demanding data centers and work as the ideal top-of-rack switches for various hyper-converged infrastructure deployments.

Equipped with intelligent ASICs, Mellanox Ethernet adapter cards provide offloading mechanisms such as Erasure Coding, T10/DIF, TCP and UDP offloads, and overlay offloads

By bypassing the CPU, RDMA over Converged Ethernet (RoCE) frees up the CPU's resources, especially in a CPU-bound environment, for the necessary storage and compute tasks, allowing for higher scalability and efficiency with the data center.


MELLANOX SPECTRUM™ SWITCHES BENEFITS:

- Cutting-edge 25/50/100G Ethernet switches with line-rate performance.
- Non-blocking, cut-through switching at line speeds of 1/10/25/40/50/100Gb/s with no packet loss.
- The dynamically shared switch buffer provides the best microburst adoption.
- Unleash the maximal power of NVMe storage pool across the network.

SOLIDSCALE REFERENCE ARCHITECTURE


SolidScale is a validated reference architecture that implements the next-generation of intelligent infrastructure platforms designed to provide applications with all of the raw performance you can expect from local storage — and all of the flexibility, manageability and scalability typically experienced with traditional storage area network-based solutions. Leveraging the latest NVMe and PCIe standards, this reference architecture uses a low-latency, high-bandwidth networking to connect compute and storage together in a flexible way that will fit almost any application's architecture requirements. Deploy this RA in storage-centric, compute-centric or mixed storage and compute configurations, and integrate it with existing data center application servers as desired.

With the SolidScale Reference Architecture, you can unlock unused capacity and performance, so you can run the same workload with potentially fewer storage devices because capacity can be shared across applications and servers. This reference architecture has the horsepower to host the most demanding, real-time analytics applications, enterprise databases, and high-performance burst buffers. The solution is available in one of two configurations: one, a block-only disaggregated storage server; two, a converged system capable of running a file server solution (such as an integrated burst buffer) or your application of choice.




SPECS

SUPERSERVER 2029BT-HNR



KEY APPLICATIONS

- Compute Intensive Application
- HPC
- Data Center
- Enterprise Server
- Hyperscale / Hyperconverged



KEY FEATURES

Four hot-pluggable systems (nodes) in a 2U form factor. Each node supports the following:

- Dual socket P (LGA 3647) supports Intel® Xeon® Scalable Processors, Dual UPI up to 10.4GT/s
- 2 PCI-E 3.0 x16 (LP) slot; 1 SIOM card support (flexible networking)
Note: must bundle with Network card
- Up to 3TB ECC 3DS LRDIMM, up to DDR4-2666MHz; 24 DIMM slots

CPU

- Dual Socket P (LGA 3647)
- Intel® Xeon® Scalable Processors, Dual UPI up to 10.4GT/s
- Support CPU TDP 70-205W* with IVR

SERVERBOARD

Super X11DPT-B

CHIPSET

Intel® C621 chipset

MEMORY

- 24 DIMM slots
- Up to 3TB ECC 3DS LRDIMM, 3TB ECC RDIMM, DDR4 up to 2666MHz

EXPANSION

2 PCI-E 3.0 x16 (LP) slot; 1 SIOM card support (flexible networking)

DRIVE BAYS

6 Hot-swap 2.5" NVMe drive bays

CONNECTIVITY

SIOM Networking

MANAGEMENT

- IPMI 2.0
- Intel® Node Manager
- KVM with dedicated LAN
- SSM, SPM, SUM
- SuperDoctor® 5
- Watch Dog

POWER SUPPLY

2200W Redundant Power Supplies with PMBus™

COOLING SYSTEM

4 Heavy duty 8cm PWM fans with air shroud

FORM FACTOR

2U Rackmount